

AMENDMENTS TO THE CLAIMS:

Please amend claims 37, 39, 44, 46, 47, 49, 54, 56, 61, 63, 64, 66, and 70 as indicated below. This listing of claims will replace all prior versions and listings of claims in the application.

Listing of Claims:

1.-35. (Cancelled)

36. (Previously Presented) A method of monitoring operation of a processing system, comprising system resources and having a plurality of processes running thereon, comprising the step of monitoring, for at least two processes in said plurality, a set of system primitives that allocate or release said system resources.

37. (Currently Amended) The method of claim 36, wherein said set of primitives monitored comprises all ~~the~~ system primitives that allocate or release said system resources.

38. (Previously Presented) The method of claim 36, wherein said set of primitives monitored comprises exclusively those system primitives that allocate or release said system resources.

39. (Currently Amended) The method of claim 36, wherein monitoring said system primitives comprises at least one of:

tracking the processes running on said system and monitoring resources used thereby[.,,];

monitoring connections by said processes running on said system[.,,];

monitoring ~~the~~ file-related operations performed within said system[.,,]; and

monitoring operation of commonly used modules with said system.

40. (Previously Presented) The method of claim 36, wherein said set of primitives monitored identifies a state of said processing system, the method further comprising the steps of:

recording a current state of said system over a current period of time and a previous state of the system over a previous period of time;

revealing any differences between said current state of the system and said previous state of the system; and

detecting any such difference revealed as a likely anomaly in the system.

41. (Previously Presented) The method of claim 40, wherein said anomaly detection comprises a learning stage to generate said previous state of the system based on said learning stage.

42. (Previously Presented) The method of claim 40, wherein said anomaly detection comprises the step of correlating a plurality of said anomalies detected and deciding whether these identify a dangerous event for the system.

43. (Previously Presented) The method of claim 42, comprising the step of emitting an alert signal indicative of any dangerous event for the system identified.

44. (Currently Amended) The method of claim 42, comprising the steps of:
generating a sequence of said anomalies;
producing a sequence of pre-conditions in a rule base; and
if said sequence of anomalies ~~at least loosely~~ matches said sequence of pre-conditions, issuing a resulting alert signal.

45. (Previously Presented) The method of claim 42, comprising the step of assigning respective weights to said anomalies in said plurality, each said weight being indicative of the criticality of the event represented by the anomaly to which the weight is assigned.

46. (Currently Amended) The method of claim 43, wherein said step of correlating comprises associating with each anomaly a value of the a weight at the previous alert signal emission time plus the current value modulated with an exponential decay factor, whereby the significance thereof decreases over time.

47. (Currently Amended) The method of claim 46, wherein said processing system operates on process identifiers, whereby a plurality of anomalies are detected for the same process identifier and said anomalies are aggregated over time according to the following formula:

$$W_{i+1}(t) = W_i(T_{i+1} - T_i) + LA_{i+1} \cdot \exp\left(-\frac{t - T_i}{\tau}\right)$$

$$W_0 = 0$$

where W_i is the weight of a user level alert signal associated with the a common stream of anomalies, when the i -th anomaly is detected; T_i is the time of detection of the i -th anomaly, LA_i is the a weight associated to the i -th anomaly, and τ is a time-decay constant.

48. (Previously Presented) The method of claim 42, wherein said step of correlating comprises the step of mapping said anomalies in said plurality into respective fuzzy sets.

49. (Currently Amended) The method of claim 40, wherein said monitoring comprises intercepting low-level data within said a system watching for changes in the state of the system, thus providing data to be analyzed in said anomaly detection.

50. (Previously Presented) The method of claim 36, comprising the step of providing a plurality of modules for performing said monitoring, said plurality of modules comprising a first set of components depending on the system being monitored and a second set of components that are independent of the system being monitored.

51. (Previously Presented) The method of claim 50, comprising the step of providing within said first set of modules at least one module selected from the group of:

a device driver for intercepting the system calls associated with said primitives in said set,

a kernel information module configured for reading information for all processes running on said monitored system, and

a system call processor configured for reading the binary data related to the system calls of said system and translating them into respective higher-level system call abstractions.

52. (Previously Presented) The method of claim 40, comprising the step of monitoring all processes running on the system monitored and all file descriptors and the socket description used by each said process to produce an instantaneous state of the system monitored.

53. (Previously Presented) An apparatus for monitoring operation of a processing system, comprising system resources and having a plurality of processes running thereon, comprising analysis modules configured for monitoring, for at least two processes in said plurality, a set of system primitives that allocate or release said system resources.

54. (Currently Amended) The apparatus of claim 53, wherein said analysis modules are configured for monitoring all the system primitives that allocate or release said system resources.

55. (Previously Presented) The apparatus of claim 53, wherein said analysis modules are configured for monitoring exclusively those system primitives that allocate or release said system resources.

56. (Currently Amended) The apparatus of claim 53, wherein said analysis modules are selected from the group of:

at least one application knowledge module tracking the processes running on said system and monitoring resources used thereby,

a network knowledge module monitoring connections by said processes running on said system,

a file-system analysis module monitoring the file-related operations performed within said system, and

a device monitoring module monitoring operation of commonly used modules with said system.

57. (Previously Presented) The apparatus of claim 53, wherein said set of primitives monitored identifies a state of said processing system, comprising a detection component configured for recording a current state of said system over a current period of time and a previous state of the system over a previous period of time, revealing any differences between said current state of the system and said previous state of the system, and detecting any such difference revealed as a likely anomaly in the system.

58. (Previously Presented) The apparatus of claim 57, wherein said detection component is configured for running a learning stage to generate said previous state of the system based on said learning stage.

59. (Previously Presented) The apparatus of claim 57, wherein said detection component is configured for correlating a plurality of said anomalies detected and deciding whether these identify a dangerous event for the system.

60. (Previously Presented) The apparatus of claim 59, wherein said detection component is configured for emitting an alert signal indicative of any dangerous event for the system identified.

61. (Currently Amended) The apparatus of claim 59, wherein said detection component is configured for:

- generating a sequence of said anomalies;
- producing a sequence of pre-conditions in a rule base; and
- if said sequence of anomalies ~~at least loosely~~ matches said sequence of pre-conditions, issuing a resulting alert signal.

62. (Previously Presented) The apparatus of claim 59, wherein said detection component is configured for assigning respective weights to said anomalies in said plurality, each said weight being indicative of the criticality of the event represented by the anomaly to which the weight is assigned.

63. (Currently Amended) The apparatus of claim 60, wherein said detection component is configured for associating with each anomaly a value of ~~the~~ a weight at the previous alert signal emission time plus the current value modulated with an exponential decay factor, whereby the significance thereof decreases over time.

64. (Currently Amended) The apparatus of claim 63, wherein said processing system operates on process identifiers (PID), whereby a plurality of anomalies are detected for the same

process identifier, and said detection component is configured for aggregating said anomalies over time according to the following formula:

$$W_{i+1}(t) = W_i(T_{i+1} - T_i) + LA_{i+1} \cdot \exp\left(-\frac{t - T_i}{\tau}\right)$$

$$W_0 = 0$$

where W_i is the weight of a user level alert signal associated with the a common stream of anomalies, when the i -th anomaly is detected; T_i is the time of detection of the i -th anomaly, LA_i is the a weight associated to the i -th anomaly, and τ is a time-decay constant.

65. (Previously Presented) The apparatus of claim 59, wherein said detection component is configured for correlating said anomalies in said plurality by mapping them into respective fuzzy sets.

66. (Currently Amended) The apparatus of claim 57, wherein said monitoring comprises an information gathering component configured for intercepting low-level data within said a system watching for changes in the state of the system, thus providing data to be analyzed in said anomaly detection.

67. (Previously Presented) The apparatus of claim 53, comprising a plurality of modules for performing said monitoring, said plurality of modules comprising a first set of components depending on the system being monitored and a second set of components that are independent of the system being monitored.

68. (Previously Presented) The apparatus of claim 67, wherein said first set of modules comprises at least one module selected from the group of:

a device driver for intercepting the system calls associated with said primitives in said set;

a kernel information module configured for reading information for all processes running on said monitored system; and

a system call processor configured for reading the binary data related to the system calls of said system and translating them into respective higher-level system call abstractions.

69. (Previously Presented) The apparatus of claim 57, comprising a current state module monitoring all processes running on the system monitored and all file descriptors and the socket description used by each said process to produce an instantaneous state of the system monitored.

70. (Currently Amended) A computer readable medium encoded with a computer program product loadable in the into a memory of at least one computer [[and]], the computer program product comprising software code portions for performing the steps of the method of claim 36.